

APPENDIX A

Trading Primer

At first glance, the process of making a trade seems to be simple. Someone wanting to buy an item meets with a potential seller, they agree on a price, and money is exchanged in return for the item. Even the most complex trading ideas begin with this concept. The mechanics may be much more complicated—perhaps the buyer and seller negotiate through a sophisticated electronic medium. Perhaps the item is actually a sophisticated financial instrument or set of instruments. Perhaps there are complications such as currency adjustments or financing costs to be considered, or perhaps the transaction is merely arranged to occur at a future point in time. Regardless, this basic meeting of buyer and seller—weighing of value against value—is the very essence and the root of all market activity.

Many books begin by saying that a trade occurs when a buyer and a seller agree on value, but this is not entirely correct. If this were so, if the parties truly agreed that the price represented the fair value of the asset, that one was equal to the other, wouldn't they each be willing to immediately unwind the trade and even to take the other side? This is almost never the case. In simple buying or selling transactions (excluding spreading and hedging, which we will get to in a minute) the buyer is willing to part with the money because he believes the asset will offer him more value in the future than the money he gave up. The seller has made a decision that the utility value of the money she will receive exceeds the value she would get from continuing to own the asset, so each participant has made an assessment of value that might be unique to his or her particular situation at that point in time. Rather than an agreement, each and every trade that occurs in the market represents a *disagreement* over the value of the money and the assets being exchanged.

THE SPREAD

We often hear language like “The stock of company XYZ is at \$50,” but even here we have an issue, for there is rarely only one price to consider; usually there are at least three. In active, liquid markets, there is a *bid* price at which buyers are willing to buy an asset, an *offer* (or *ask*, from *asking price*) at which sellers are offering to sell, and usually a *last print* (or simply, *last*) price where a trade was made. In a typical market, bids will be lined up below the market (more buyers are willing to buy at lower prices), and offers will be stacked at higher prices. The *inside market* refers to the highest bid price and the lowest offer, representing the best available price at any time, and the distance between those two prices is often referred to as the *spread*.

In the hypothetical case of XYZ that just traded at \$50, we might find that the best available bid is \$49.95 and the lowest offered price is \$50.05. A trader would verbalize this situation as “49.95 bid, offered at 50.05,” which is usually shortened in practice to “49.95 at 50.05.” If the trader is reasonably sure the person he is speaking to knows the approximate price, he may drop the *handle* (the whole number) part of the price and just give the decimal pricing. As prices change rapidly, this can lead to a dialogue like “95 at 05 [meaning .95 and .05], at 7, at 9, 98 bid, now at 06.” Though the last trade may have been at \$50, there is a reasonable chance that we could not execute at that price in this situation with the market “95 at 05.” In fact, it is even possible that the last print was 50.00, and now the market may have moved to 50.50 at 50.60. A buyer, in this case, has little hope of executing at any price under 50.50; the spread is often a better reflection of actual value than the last print.

Some traders glean a lot of information from the spread. For instance, the size of the spread is often a measure of the uncertainty in the market; when prices are changing rapidly, market participants often react by widening the spread. Buyers are not aggressive in paying high prices because they know the market could move against them in the next instant, so bids tend to drop lower. At the same time, sellers also react by lifting their offers to higher prices. There is other information in the spread: how rapidly it moves, exactly how it moves (do sellers lift offers higher or do buyers aggressively bid higher?), how much size is displayed, and many other subtleties.

If this seems like a lot of time spent on something very basic, you’re right; it is, and it is also very important. The spread represents a very real cost of trading. Imagine that a buyer pays the offer, and then immediately realizes he has made a mistake and wants to get out of the trade. This can be done only by turning around and selling to the buyers on the bid, so the spread is a source of risk and a cost of trading in every transaction. It is also usually the best estimation of the actual market value of an asset. In the case of a very active stock that prints a trade every few seconds, the inside market is rarely very far from the last print, so the last trade is a very good approximation of value. (This kind of stock will usually tend to have a relatively tight spread as well.)

However, there are some instruments that may go days without trading (some stocks, options, distant months of futures contracts, for instance). In these cases, the last

price may be completely irrelevant because it happened so long ago that the market has moved. To further complicate matters, very inactive, illiquid instruments will often have extremely wide spreads. If stock XYZ last printed 50.00 three weeks ago, but is now bid at 30.00 and offered at 49.00, what is it actually worth? Researchers studying price patterns need to be careful because printed price records in illiquid markets can be very misleading. As an interesting aside, a major factor in the 2007–2008 financial crisis was the importance of many financial instruments representing very significant financial commitments that did not have liquid markets. Spreads were wide, or in some cases, nonexistent, so it was impossible to derive a market value for many of these assets. In the absence of a market price, traders resorted to building complex models with many moving parts (if that sounds a lot like guessing, you're not wrong), and many of these models gave very misleading values for these instruments. This complete breakdown in understanding the value and risk of these instruments was one of the major contributors to the crisis.

Liquidity is a misused and often imprecisely defined term, but it usually means the availability of willing buyers and sellers. Liquid instruments tend to have tighter spreads and *deeper books*, meaning that there are many buyers and sellers at price levels beyond the inside market. Going back to XYZ, which is now 49.95 at 50.05, we might find there are 1,000 shares on the bid at 49.95, and many thousands at 49.94, 49.93, and so on for many pennies below the market. Imagine a very large sell order comes into the market. These buyers would easily be able to absorb that order, meaning that XYZ would trade on the bid at 49.95, then maybe 49.94 and 49.93—all in all, very little price change. However, imagine a second scenario where there are again 1,000 shares on the bid at 49.95, but now a few hundred at 49.91, a few more at 49.87, and so on. If a large sell order hits this market, it will “clear the bids” and the price will drop much lower.

Market makers are a specific group of traders whose job is basically to provide liquidity. A market maker will usually have both a bid and an offer in the market, though they are free to adjust those levels as needed. For instance, if market makers are getting hit on the bid so that they are accumulating large long positions, they may choose to still bid for the stock, but to drop their bid to lower levels. If they keep their offer price on the inside market (meaning that they adjust their offer so that it is the lowest offer), eventually they will be able to sell some of their inventory, and in this way manage their exposure. Floor traders in open-outcry markets were the original market makers. Early electronic markets had designated market makers, but this role has now passed to many firms who run computer programs (algorithms) that function as market makers.

Market makers incur significant risks at times because they will always be on the wrong side of big moves driven by informed traders. If a group of traders comes into the market with many buy orders, the market maker will be forced to short to take the other side of that trade. In extreme cases (e.g., the crash of 1987), market makers could be forced out of business by adverse price moves. Market makers are typically offered various incentives to compensate them for these risks; otherwise, no one would take this job! (As of this writing, in 2011, liquidity rebates, which pay the firm a very small fee for executing via limit orders, are one of the primary forms of compensation for most traders functioning as market makers.)

TWO TYPES OF ORDERS

Traders wishing to buy this market have, broadly speaking, two options. A buyer who is not really in any hurry might focus on getting the best (lowest) price he can get, so he can *bid* for the stock. (In this case, bid is used as a verb meaning “to place an order lower than current prices.”) Consider the case of XYZ with the market bid at 49.95 and offered at 50.05; a buyer could *join the bid* and put his order at 49.95. Of course, the buyer can also place orders lower, but they will be filled only if the market moves down to the level where he is bidding. If the buyer is feeling a little more urgency, he can *step in front of the bid* and put his order at 49.96 or 49.97. Note that in this case, the inside market would now be 49.97 (still offered at 50.05). This is a natural force that tends to compress spreads as buyers bid slightly higher and sellers offer lower in competition for fills, and is one of the main reasons why active, liquid markets tend to have tight spreads. However, if the buyer is really motivated and must have stock XYZ now, he may choose to *pay the offer* (other common language is *take the offer*, and the reverse is to *hit the bid*). Though much simplified, these are the two options available to traders and they correspond to the two most commonly used order types.

Limit orders are orders where buyers try to buy at a cheaper price than the offer and sellers offer to sell at prices above the bid. A trader wanting to buy XYZ with a limit order might say, “Bid for it” or “Join the bid,” as in “We’re in no hurry here—just bid for it.” A seller might say, “Okay, fine, offer it out. Put it up on the offer.” *Market orders* are orders that will execute immediately. Buyers will pay the offer and sellers will end up hitting the bid. These orders usually reflect some degree of urgency—the order must be done immediately and a better price (trade location) is sacrificed for speed of execution. Though the buyer bidding with a limit order will get a better price if filled, the trade-off is that the order may never be filled if the market moves higher.

Now we have arrived at something subtle and very important. Imagine XYZ is again 49.95 at 50.05, and then it trades at 50.05. What just happened? A buyer wanted to buy the stock so much that he was willing to *pay the spread* or take the offer; we can say that this is a buyer-motivated trade. By analyzing the number of orders that hit the bid (seller-motivated trades) compared to those that take the offer (buyer-motivated trades), whether through computer-aided analysis or careful observation, traders can get a deeper sense of the conviction levels and urgency behind price moves. For instance, a stock may move from 50 to 51 in a series of back-and-forth motions with orders printing on both the bid and the offer. At another time, the same price change could occur in a straight line as buyers keep paying the offer and keep that buying pressure on the offer for the entire move. Simple observation would note that both moves began at 50 and ended at 51, and perhaps even occurred in the same amount of time and with the same amount of volume being done, but each of these moves suggests something different about the underlying conviction in the market.

A trader who buys something in anticipation of it going higher is said to be *long* that instrument. For most people, this is a natural and intuitive concept: buying something relatively cheaply, planning to sell it later for a higher price, and pocketing the difference (minus any costs of financing or insurance incurred in the interim). One of the divisions

between professional and amateur traders is that professionals are often just as willing to *go short* (or just to *short*) a market, but the public often has a bias against shorting. There are several reasons for this; most equity traders have a natural inclination toward owning stocks, and think that shorting is a very complex transaction, or that it is somehow immoral to bet on a company's value going down. (In some markets it is theoretically complex, as the instrument must first be borrowed, then sold, later bought back, and finally returned to the lender.) This prejudice is unfounded and is one of the key differences between the public and most professionals. Shorting is nothing more than the opposite of being long. Whereas the buyer seeks to make a profit as prices rise, the short seller anticipates falling prices and hopes to *cover the short* (buy it back) at lower prices. Short selling is an important part of the trader's tool set.

Spread Trading

It is also worth considering that a certain amount of the buying and selling pressure in the market represents more complex interests than simple buying and shorting. Imagine a farmer who knows he will have a grain crop coming in September. This farmer might sell his grain in the futures market before the harvest comes in (technically, a short sale), but this does not in any way mean that he expects that prices will be going down. This is a simple example of a *hedging* transaction, and more complicated examples exist in all markets.

Spread trading is another type of transaction involving buying one asset and selling another, looking to profit from the change in value between the two. Imagine that a trader feels that Assets A and B should trade in a more or less predictable relationship, perhaps with Asset A at a premium due to production costs and so forth. This trader could track the spread, or the difference between Assets A and B. If she feels that the premium is too small, she can go long the spread by buying Asset A and shorting Asset B. This is not a bullish or bearish bet on either of the assets, but a bet that the spread between the two will widen. She will make money if A goes up and B goes down, if they both go up but A goes up more, or even if they both decline as long as B declines more. The key is that the percentage change of A must be greater than the percentage change of B (assuming an equal-weighted position) from the time the trade is entered. Conversely, if the trader felt the spread was too wide, she could short the spread by doing the reverse of this transaction.

This just scratches the surface of these complex, multileg transactions, which exist in and between all markets in virtually unlimited combinations. The message here is that much of the buying and selling we observe in markets may be part of spread trades like this. In this case, we might see the trader buying Asset A and assume she was bullish on it, when in fact, she does not care if it goes up or down—all she cares about is the spread relationship. Spreads can be created between different asset classes, regional markets, or international markets, and these trades can even be initiated with different timings on each leg. Do not assume that buying is always bullish and selling is always bearish; there may be much more going on behind the scenes.

CHARTS

The most basic language of the market is price changes, perhaps with associated volume. In the case of XYZ, we might see 500 shares done at 50.00, 300 @ 50.03, 500 @ 50.04, and so on. (A more complete record would include a time stamp and whether the order was executed on the bid, on the offer, or in between.) This record is sometimes referred to as the “prints on the tape” in memory of the old-style ticker tape machines, which really did print prices on paper tape. A very active instrument can “print” hundreds or thousands of trades in a single minute, so a trader would quickly become lost without some reference of historical activity. A chart of these price changes is the natural solution, and many discretionary traders find that charts present market information in a way that is both intuitive and useful.

Tick Charts

The most primitive type of chart would simply be the information in the prints plotted on a graph with price on the y-axis and trade number on the x-axis. The chart in Figure A.1




FIGURE A.1 A Single-Tick Chart of AAPL

is a *single-tick* chart of the market’s opening for Apple, Inc. (Nasdaq: AAPL) on 7/12/2011 (trade numbers not shown). The single-tick chart is interesting from an intellectual perspective because it is the most fundamental language of the market—every single trade, every transaction that hits the *consolidated tape*, is displayed on the chart. From a practical perspective, a single-tick chart is not very useful, for the simple reason that it is difficult to fit much trading history on a single chart. This AAPL chart shows only the first 30 seconds of the trading day. In an extremely active stock, it is even possible that less than a single second’s trading could fit in the same space; it is not possible for the human eye and brain to process this information in a meaningful manner.

The obvious solution is to aggregate many ticks into a single space on the chart. When this is done, we refer to the chart by how many ticks (trades) are put together into a single space on the x-axis. In Figure A.2, each bar represents 25 individual trades, and, now, approximately five and a half minutes of activity fit on the same chart space. It is important to remember that the x-axis is not scaled to time. This is one of the main advantages of aggregate tick charts: as the market becomes more or less active the x-axis expands or contracts to accommodate the activity. In many cases, this can



FIGURE A.2 A 25-Tick Chart of AAPL




FIGURE A.3 Comparison of Bar and Candlestick Charts

create some advantages over simple time-scaled charts by making readable patterns out of very illiquid or extremely volatile markets.

Bars and Candles

Once we start aggregating trades on the x-axis, a graphical device is needed to explain the activity in that space. One logical solution is simply to plot a vertical bar, with the top and bottom of the bar representing the high and low extremes reached during the period (right panel in Figure A.3). Traders often find it useful to mark the first trade of the period with a tick on the left side of the bar (the *open*) and the last trade of the period with a tick on the right side (the *close*). These four data points, plotted like this, create the classical bar chart. An alternate format is the *candlestick chart*, which prints a wide body between the open and close of the bars, with thinner *shadows* (also called “tails” or “wicks”) above and below the body, reaching to the high and low of the period (left panel in Figure A.3). Candlestick charts were first used in medieval Japan for rice trading, so much of the terminology associated with patterns of candles is also Japanese. Traditionally, the body of the candle is filled in if the close is lower than the open and left empty if the close is

higher than the open. Modern charting software packages usually fill the body of every candle but change the color to something intuitive like red for downward-closing candles and green or blue for upward-closing candles.

Time Charts and Other Options

Though the strength of tick charts is their ability to adapt to activity levels in the markets, traders and analysts often prefer a format where each bar corresponds to a predictable unit of time. These time charts are by far the most commonly used in most applications, and are especially useful for traders looking at daily and longer time frames. The chart is referred to by the length of each time unit (as in a 5-minute chart), which may also be called the *time frame* of the chart. Many traders choose to look at different time frames (for instance, 5-minute, hourly, and daily charts) for the same market to get a better sense of the forces affecting prices at any time.

There are other possibilities for scaling the x-axis, but these tend to be less used except in certain specific contexts. It is possible to aggregate bars not by ticks (transactions) but by trading volume. In this case, a bar would contain a certain number of




FIGURE A.4 Market Profile Chart of AAPL

shares or contracts, after which the next bar would begin. Another possibility is to have each bar end after a certain range is reached, so that, for instance, each bar would be 0.25 wide. (Be aware that the range bars created at the end of a trading day will be different from the ones created as the market unfolds. This makes backtesting and analysis on range bars virtually impossible.)

Other types of charts include swing charts, point and figure charts, kagi charts, line charts, and many others, but these are much less common and are outside the scope of this primer, with one exception. One other chart worth mentioning is Market Profile, created by Peter Steidlmayer in the 1980s. Market Profile charts essentially allow the trader to look inside the bar to see where most volume and trading activity occurred in the time period. Market Profile displays this information in a graphical format attached to the right of each bar, with wider horizontal bars indicating more trading activity at that price level. (See Figure A.4.) While most charting methods only reprocess the same open, high, low, and close (OHLC), volume, and open interest information, Market Profile is an important innovation—it adds detail and perspective that is not visible on a standard chart.